

CLAIMS

- 1 1. A method for distributing parity across a disk array, the method comprising the steps
2 of:
 - 3 adding a new disk to pre-existing disks of the array;
 - 4 dividing each disk into blocks, the blocks being organized into stripes such that
5 each stripe contains one block from each disk; and
 - 6 distributing parity among blocks of the new and pre-existing disks without recal-
7 culation or moving of any blocks containing data.
- 1 2. The method of Claim 1 wherein the step of distributing comprises the step of distrib-
2 uting parity among blocks of the new and pre-existing disks in a manner that maintains a
3 fixed pattern of parity blocks among stripes of the disks.
- 1 3. The method of Claim 1 wherein the step of distributing comprises the step of changing
2 an assignment for one or more blocks containing parity of each pre-existing disk to the
3 newly added disk.
- 1 4. The method of Claim 2 wherein the step of adding comprises the step of initializing
2 the added disk so as to not affect parity of the stripes.
- 1 5. The method of Claim 4 wherein the step of initializing comprises the step of reas-
2 signing blocks containing parity in certain stripes to the new disk without calculation or
3 writing of parity.
- 1 6. The method of Claim 5 wherein the certain stripes comprise $1/N$ of the stripes, where
2 N is equal to the number of disks in the array.

- 1 7. The method of Claim 5 wherein the step of reassigning comprises the step of chang-
2 ing a block containing parity (parity block) to a block containing data (data block) and
3 not changing a data block to a parity block.
- 1 8. The method of Claim 1 wherein the step of distributing comprises the step of reas-
2 signing one of N blocks containing parity (parity blocks) from each pre-existing disk to
3 the added disk, wherein N is equal to the number of disks in the array.
- 1 9. The method of Claim 8 wherein the step of reassigning comprises the step of reas-
2 signing one of N parity blocks to the new disk, with each pre-existing disk continuing to
3 hold 1/N of the parity blocks in the array.
- 1 10. A system adapted to distribute parity across disks of a storage system, the system
2 comprising:
3 a disk array comprising pre-existing disks and at least one new disk; and
4 a storage module configured to compute parity in blocks of stripes across the
5 disks and reconstruct blocks of disks lost as a result of failure, the storage module further
6 configured to assign the parity among the blocks of the new and pre-existing disks with-
7 out recalculation or moving of any data blocks.
- 1 11. The system of Claim 10 further comprising a table configured to store parity assign-
2 ments calculated for one of a known group size of the disk array and a maximum group
3 size of the array, the stored parity assignments defining a repeat interval of a parity distri-
4 bution pattern used to determine locations of parity storage on any disk in the array.
- 1 12. The system of Claim 10 wherein the storage module is embodied as a RAID system
2 of the storage system.

- 1 13. The system of Claim 10 wherein the storage module is embodied as an internal disk
2 array controller of the storage system.
- 1 14. The system of Claim 10 wherein the storage module is embodied as a disk array
2 control system externally coupled to the storage system.
- 1 15. The system of Claim 10 wherein the disk array is a block-based RAID array.
- 1 16. A method for distributing commodities over containers of a system, the method
2 comprising the steps of:
3 adding a new container to pre-existing containers of the system to thereby provide
4 N containers; and
5 moving only $1/N$ of the commodities to the new container.
- 1 17. The method of Claim 16 wherein the system is a storage system, the commodities are
2 data structures adapted for storage on storage devices of an array, and the containers are
3 storage entities coupled to the array.
- 1 18. The method of Claim 17 wherein the storage entities are storage heads.
- 1 19. The method of Claim 17 wherein the data structures are inode file blocks.
- 1 20. Apparatus for distributing parity across a disk array, the apparatus comprising:
2 means for adding a new disk to pre-existing disks of the array;
3 means for dividing each disk into blocks, the blocks being organized into stripes
4 such that each stripe contains one block from each disk; and
5 means for distributing parity among blocks of the new and pre-existing disks
6 without recalculation or moving of any blocks containing data.

- 1 21. A computer readable medium containing executable program instructions for distrib-
- 2 uting parity across a disk array, the executable instructions comprising one or more pro-
- 3 gram instructions for:
 - 4 adding a new disk to pre-existing disks of the array;
 - 5 dividing each disk into blocks, the blocks being organized into stripes such that
 - 6 each stripe contains one block from each disk; and
 - 7 distributing parity among blocks of the new and pre-existing disks without recal-
 - 8 culation or moving of any blocks containing data.